

Mining frequent items in the time fading model

Massimo Cafaro^{a,*}, Marco Pulimeno^a, Italo Epicoco^a, Giovanni Aloisio^a

^aUniversity of Salento, Lecce, Italy

Abstract

We present FDCMSS, a new sketch-based algorithm for mining frequent items in data streams. The algorithm cleverly combines key ideas borrowed from forward decay, the Count-Min and the Space Saving algorithms. It works in the time fading model, mining data streams according to the cash register model. We formally prove its correctness and show, through extensive experimental results, that our algorithm outperforms λ -HCount, a recently developed algorithm, with regard to speed, space used, precision attained and error committed on both synthetic and real datasets.

Keywords: frequent items, time fading model, cash register model, sketches.

1. Introduction

A data stream σ consists of a sequence of n items drawn from a universe \mathcal{U} . Without loss of generality, let m be the number of distinct items in σ i.e., let $\mathcal{U} = \{1, 2, \dots, m\}$. Depending on the applications, items may be numbers, IP addresses, points, graph edges etc. Owing to the huge size of σ , an algorithm in charge of processing its items is subject to the stringent requirement that no more than one pass over the data is allowed. In practice, storing the items is not a feasible option. We refer the interested reader to [30], a very good survey of streaming algorithms, for additional details and underlying reasons motivating research in this area. In this paper, we deal with mining of frequent items in a data stream. This problem has been extensively studied, and is recognized as one of the most important in the streaming algorithms literature, where it is also called, depending on the specific context, *hot list analysis* [21], *market basket analysis* [3] and *iceberg query* [19], [2].

Among the many possible applications, we recall here network traffic analysis [16], [18], [31], analysis of web logs [6], Computational and theoretical Linguistics [20].

Letting f_i denote the frequency of the item $i \in \mathcal{U}$ (i.e., its number of occurrences in σ), $\mathbf{f} = (f_1, \dots, f_m)$ the frequency vector, $0 < \phi < 1$ a support threshold and $\|\mathbf{f}\|_1$ the 1-norm of \mathbf{f} (which represents the total number of occurrences of all of the stream items), an approximate solution requires returning all of the items which are frequent, i.e., those items i such that $f_i > \phi \|\mathbf{f}\|_1$ and, letting $0 < \epsilon < 1$ denote the error committed, an algorithm must not return any item i such that $f_i \leq (\phi - \epsilon) \|\mathbf{f}\|_1$. In particular, the error is such that $\epsilon < \phi$.

Beyond the traditional distinction between deterministic and randomized algorithms, in this area algorithms for detecting frequent items are also often referred to as being either *counter* or *sketch* based. In counter-based algorithms, a fixed number of counters is used to keep track of stream items. Indeed, given a support threshold $0 < \phi < 1$, the number of possible frequent items is an integer belonging to the open interval $(0, 1/\phi)$. Sketch-based algorithms monitor the data stream by using a set of counters, stored in a sketch data structure, usually a bi-dimensional array. Stream items are mapped by hash functions to their corresponding cells in the sketch. Whilst counter-based algorithms are deterministic, sketch-based ones are randomized and provide a probabilistic guarantee.

The streaming model we have described so far is called *cash register* or *strict turnstile* model [30], since only *insertions* are allowed. On the other hand, in the more general *turnstile* model, *deletions* are also allowed. An

*Corresponding author

Email addresses: massimo.cafaro@unisalento.it (Massimo Cafaro), marco.pulimeno@unisalento.it (Marco Pulimeno), italo.epicoco@unisalento.it (Italo Epicoco), giovanni.aloisio@unisalento.it (Giovanni Aloisio)

advantage of sketch-based algorithms is that they can easily support deletions and can therefore work in the turnstile model; counter-based algorithms only work in the cash register model.

Among the counter based algorithms, the first sequential algorithm has been proposed by Misra and Gries [29]. Later, this algorithm was rediscovered independently by Demaine et al. [16] (the so-called *Frequent* algorithm) and Karp et al. [24]. Recently developed counters-based algorithms include *Sticky Sampling* and *Lossy Counting* [27], and *Space Saving* [28]. Notable sketch-based algorithms are *CountSketch* [6], *Group Test* [12], *Count-Min* [11] and *hCount* [23].

Regarding parallel algorithms, [5] and [4] present message-passing based parallel versions of the Frequent and Space Saving algorithms. Among the algorithms for shared-memory architectures we recall here a parallel version of Frequent [35], a parallel version of Lossy Counting [34], and parallel versions of Space Saving [32] [14]. Novel shared-memory parallel algorithms for frequent items were recently proposed in [33]. Accelerator based algorithms for frequent items exploiting a GPU (Graphics Processing Unit) include [22] and [17].

In this paper, we are concerned with the problem of detecting frequent items in a stream with the additional constraint that recent items must be weighted more than former items. The underlying assumption is that, in some applications, recent data is certainly more useful and valuable than older, stale data. Therefore, each item in the stream has an associated timestamp that will be used to determine its weight. In practice, instead of estimating frequency counts, we are required to estimate *decayed counts*. Two different models have been proposed in the literature: the *sliding window* and the *time fading* model.

In the sliding window model [15] [30], freshness of recent items is captured by a time window, i.e., a temporal interval of fixed size in which only the most recent N items are taken into account; detection of frequent items is strictly related to those items falling in the window. The items in the stream become stale over time, since the window periodically slides forward.

The time fading model [26] [8] [7] does not use a window sliding over time; freshness of more recent items is instead emphasized by *fading* the frequency count of older items. This is achieved by using a decaying factor $0 < \lambda < 1$ to compute an item's *decayed count* (also called *decayed frequency*) through decay functions that assign greater weight to more recent elements. The older an item, the lower its decayed count is: in the case of exponential decay, the weight of an item occurred n time units in the past is $e^{-\lambda n}$, which is an exponentially decreasing quantity.

This paper is organized as follows. We recall in Section 2 key definitions and concepts that will be used in the rest of the manuscript, and introduce in Section 3 the λ -HCount algorithm [7], a recently published sketch-based algorithm that detects frequent items in the time fading model. Then, we introduce in Section 4 our FDCMSS algorithm and formally prove in Sections 5 and 6, respectively, its error bound and correctness. Next, we compare λ -HCount to FDCMSS from a theoretical perspective, and show that our algorithm achieves its error bound using a tiny fraction of the space required by λ -HCount. Then, we provide extensive experimental results in Section 8, in which we compare again FDCMSS versus λ -HCount from a quantitative, practical perspective, and show that FDCMSS outperforms λ -HCount with regard to speed, space used, precision attained and error committed on both synthetic and real datasets. Finally, we draw our conclusions in Section 9.

2. Key ideas

Our algorithm cleverly combines ideas borrowed from forward decay [9], the Count-Min sketch-based algorithm [11], and the Space Saving counter-based algorithm [28]. In this section, we recall preliminary definitions and key concepts.

Definition 1. Given an item i with arrival time t_i , a *decay function* returns a weight for the item. In our algorithm the weight $w(i, t)$ determined at time t , depends on the timestamp t_i associated to the item. Decay functions satisfy the following properties: (i) $w(i, t) = 1$ when $t_i = t$ and $0 \leq w(i, t) \leq 1$ for all $t \geq t_i$; (ii) w is monotone non-increasing as time increases, i.e., $t' \geq t \implies w(i, t') \leq w(i, t)$.

Regarding the time fading model, related work has mostly exploited *backward decay* functions, in which the weight of an item is a function of its age, a , where the age at time $t > t_i$ is simply $a = t - t_i$. The term backward decay stems from the aim of measuring from the current time back to the item's timestamp.

Definition 2. A backward decay function is defined by a positive monotone non-increasing function f so that the weight of the i th item with arrival time t_i determined at time t is given by $w(i, t) = \frac{f(t-t_i)}{f(t-t)} = \frac{f(t-t_i)}{f(0)}$. The denominator in the expression normalizes the weight, so that it obeys condition (i) of Definition 1.

Prior algorithms and applications have been using backward exponential decay functions such as $f(a) = e^{-\lambda a}$ with $\lambda > 0$. In our algorithm, we use instead a forward decay function, defined as follows. Under forward decay, the weight of an item is computed on the amount of time between the arrival of an item and a fixed point L , called the *landmark* time, which, by convention, is some time earlier than the timestamps of all of the items. The idea is to look forward in time from the landmark to see an item, instead of looking backward from the current time.

Definition 3. Given a positive monotone non-decreasing function g , and a landmark time L , the forward decayed weight of an item i with arrival time $t_i > L$ measured at time $t \geq t_i$ is given by $w(i, t) = \frac{g(t_i-L)}{g(t-L)}$.

When $t = t_i$ the weight is 1 (condition (i) of Definition 1). Since g is monotone non-decreasing, as t increases the weight does not increase, and $0 \leq w(i, t) \leq 1$.

It's easy to prove that backward and forward exponential decay coincide. In our algorithm, we could use the exponential decay; however, we prefer to use the polynomial forward decay function $g(n) = n^2$. It has been proved that this class of functions (i.e., $g(n) = n^\beta$) satisfies a *relative decay* property, which states that for any time t after a landmark time L , the weight for items with timestamp $\gamma t + (1 - \gamma)L$ is the same. In practice, relative decay holds for forward decay functions assigning the weight of an item depending only on where the item falls as a fraction in the window defined by L and t . As an example, a function for which relative decay holds assigns to an item arriving half way between L and t the same weight as t increases.

Intuitively, this property requires assigning to an item a weight which is a function of its relative age, i.e., its age as a fraction of the total time period observed. Since backward decay is only concerned with absolute age, it can not provide relative decay. An important consequence of relative decay is that it allows selecting a meaningful landmark time L to choose for forward decay.

Besides its flexibility (e.g., choosing an appropriate polynomial function we can select and control a slower rate of decay with regard to an exponential), another advantage of forward decay is related to its ability to deal with out of order arrival of stream items. Indeed, forward decay does not rely on items arriving in increasing order of timestamps. On the contrary, under backward decay, handling out of order arrivals can require significant effort to accommodate.

Definition 4. The *decayed count*, C , of a stream σ of n items is the sum of decayed weights of items: $C = \sum_{i=1}^n \frac{g(t_i-L)}{g(t-L)}$.

We can now formally state the problem solved by our algorithm: *approximate frequent items under forward decay*.

Definition 5. (Frequent items under forward decay) For each item in the input, v , its decayed count is given by $f_v = \sum_{v_i=v} \frac{g(t_i-L)}{g(t-L)}$. Given a threshold value ϕ , the frequent items are all of the items v satisfying $f_v > \phi C$.

Definition 6. (Approximate frequent items under forward decay problem) Given an error bound ϵ and a threshold ϕ , determine all of the items satisfying $f_v > \phi C$, and report no items with $f_v \leq (\phi - \epsilon)C$.

Our goal is to design an algorithm solving the *Approximate frequent items under forward decay* problem by providing the following (ϵ, δ) approximation.

Definition 7. $((\epsilon, \delta)$ approximation) Let $A(\sigma)$ denote the output of a randomized streaming algorithm A on input σ ; it is worth noting here that $A(\sigma)$ is a random variable. Moreover, let $f(\sigma)$ be the function that A is supposed to compute. The algorithm A (ϵ, δ) approximates f if $\Pr[|A(\sigma) - f(\sigma)| > \epsilon] \leq \delta$.

Count-Min is based on a sketch whose dimensions are derived by the input parameters ϵ , the error, and δ , the probability of failure. In particular, for Count-Min $d = \lceil \ln 1/\delta \rceil$ is the number of rows in the sketch and $w = \lceil e/\epsilon \rceil$ is the number of columns. Every cell in the sketch is a counter, which is updated by hash functions. By using this data structure, the algorithm solves with high probability (i.e., with probability greater than or equal to $1 - \delta$) the *frequency estimation* problem for arbitrary items. The algorithm may also be extended to solve the *approximate frequent items* problem as well, by using an additional heap data structure which is updated each time a cell is updated. Since in

Count-Min the frequencies stored in the cells overestimate the true frequencies, a point query for an arbitrary item simply inspects all of the d cells in which the item is mapped to by the corresponding hash functions and returns the minimum of those d counters.

In addition, Count-Min allows us reusing the same underlying data structure to solve (if needed), beside frequent items, additional problems related to the same input stream (e.g., quantiles, frequency estimation, medians, etc). Moreover, the Count-Min data structure requires less space with regard to other sketch-based algorithms.

We shall prove later (see Theorem 2) that in our algorithm, with high probability, if an item i is frequent, then it appears as a majority item candidate in at least one of the d sketch cells in which it falls. Therefore, in order to detect frequent items, we decided to use the Space Saving algorithm [28]. This is a counter-based algorithm, designed to solve the *approximate frequent items* problem using $k \geq \lceil 1/\phi \rceil$ counters in order to determine the frequent items in the input data stream.

We could use a different counter-based algorithm (e.g., *Frequent* [16]); our choice stems from the well-known fact that Space Saving provides the greatest accuracy (precision, total and average relative error) among the counter-based algorithms [10] [25]. Moreover, exploiting in FDCMSS the Space Saving algorithm within the sketch cells allows us avoiding the need for an additional, separate data structure to keep track of frequent items. We now briefly recall how Space Saving works.

Let S denote the Space Saving stream summary data structure. Updating S upon arrival of an item works as shown in the pseudocode of Algorithm 1. We denote by $c_j.i$ and $c_j.f$ respectively the item monitored by the j th counter of S and its corresponding estimated frequency. When processing an item which is already monitored by a counter, its estimated frequency is incremented by its weight w . When processing an item which is not already monitored by one of the available counters, there are two possibilities. If a counter is available, it will be in charge of monitoring the item and its estimated frequency is set to its weight w . Otherwise, if all of the counters are already occupied (their frequencies are different from zero), the counter storing the item with minimum frequency is incremented by its weight w . Then the monitored item is evicted from the counter and replaced by the new item. This happens since an item which is not monitored can not have a frequency greater than the minimal frequency. The complexity of the Space Saving update procedure is $O(1)$.

Let σ be the input stream and S the stream summary data structure at the end of the sequential Space Saving algorithm's execution. Moreover, let $\sum_{c_j \in S} c_j.f$ be the sum of the counters in S , f_v the exact frequency of an item v , \hat{f}_v its estimated frequency, $\mathbf{f} = (f_1, \dots, f_m)$ the frequency vector, \hat{f}^{min} the minimum frequency in S and $\hat{\epsilon}_v$ the estimated error of item v , i.e. an over-estimation of the difference between the estimated and exact frequency.

Finally, denote by S_ϕ the set of counters in S which are monitoring items ($|S_\phi| \leq k$). It is worth noting here that $\hat{f}^{min} = 0$ when $|S_\phi| < k$. The following relations hold (as shown in [28]):

$$\sum_{c_j \in S} c_j.f = \|\mathbf{f}\|_1, \quad (1)$$

$$\hat{f}_v - \hat{f}^{min} \leq \hat{f}_v - \hat{\epsilon}_v \leq f_v \leq \hat{f}_v, \quad v \in S_\phi, \quad (2)$$

$$f_v \leq \hat{f}^{min}, \quad v \notin S_\phi, \quad (3)$$

$$\hat{f}^{min} \leq \left\lfloor \frac{\|\mathbf{f}\|_1}{k} \right\rfloor. \quad (4)$$

Therefore, it holds that

$$\hat{f}_v - f_v \leq \hat{f}^{min} \leq \left\lfloor \frac{\|\mathbf{f}\|_1}{k} \right\rfloor, \quad v \in \mathcal{U}. \quad (5)$$

To recap, we end this Section summarizing the reasons for combining forward decay, the Count-Min and the Space Saving algorithms:

- forward exponential decay coincides with backward exponential decay, so that we can still use exponentials;

Algorithm 1 Space Saving update

Require: S , a stream summary; j , an item; w , the weight of item j

Ensure: a stream summary S containing frequent items

```
1: procedure SPACEAVINGUPDATE( $S, j, w$ )
2:   if  $j$  is monitored then
3:     let  $c_l$  be the counter monitoring  $j$ 
4:      $c_l.f \leftarrow c_l.f + w$ 
5:   else
6:     if there is a counter  $c_r$  which is not monitoring any item then
7:        $c_r.i \leftarrow j$ 
8:        $c_r.f \leftarrow w$ 
9:     else
10:      let  $c_s$  be the counter monitoring the item with least hits
11:       $c_s.i \leftarrow j$ 
12:       $c_s.f \leftarrow c_s.f + w$ 
13:     end if
14:   end if
15: end procedure
```

- forward decay allows for greater flexibility (e.g., choosing a polynomial function we can select a different, slower rate of decay with regard to an exponential);
- forward decay allows easily dealing with out of order arrival of stream items;
- forward decay satisfies a relative decay property, which states that for any time t after a landmark time L , the weight for items with timestamp $\gamma t + (1 - \gamma)L$ is the same;
- the Count-Min algorithm allows us reusing the same underlying data structure to solve (if needed), beside frequent items, additional problems related to the same input stream (e.g., quantiles, frequency estimation, medians, etc);
- the Count-Min data structure requires less space with regard to other sketch-based algorithms;
- we use Space Saving to detect frequent items within the sketch cells;
- the Space Saving algorithm provides the greatest accuracy (precision, total and average relative error) among the counter-based algorithms;
- the Space Saving algorithm allows us avoiding the need for an additional, separate data structure to keep track of frequent items.

3. The λ -HCount algorithm

We now introduce λ -HCount [7], a recently published sketch-based algorithm that detects frequent items in the time fading model by using a backward decay exponential function. The λ -HCount algorithm, shown in pseudocode as Algorithm 2, requires a two dimensional sketch D of size $r \times m$ to store decayed weights and timestamps, and a doubly linked list F to store frequent items candidates, accessed through a hash function. λ -HCount is based on the use of r FNV hash functions $h_i(x), i = 1, \dots, r$ which uniformly and independently map an item to an integer in the interval $[1, m]$; the algorithm requires a support threshold s and an error bound ϵ . The occurrence of an item at time t_a is weighted in time by a factor λ^{t-t_a} , where λ represents the fading factor ($0 < \lambda < 1$). The decayed count of an item is given by the sum of its decayed weight in time. The decayed count of the stream σ , as proved by the authors, is bounded by $\frac{1}{1-\lambda}$.

Each entry $D[i, h_i(x)]$ in the sketch D stores $D[i, h_i(x)].s$ which is the decayed count of item x (also called *density*), and $D[i, h_i(x)].t$ which is the last time the value of $D[i, h_i(x)].s$ was updated. Whenever an item x arrives, its decayed count is updated in all of the r cells $D[i, h_i(x)]$, $i = 1, \dots, r$. Then, the algorithm computes $\hat{f}_x = \min_{1 \leq i \leq r} \{D[i, h_i(x)].s\}$ as the estimated decayed count of x ; if \hat{f}_x is greater than or equal to the threshold $\frac{s-\epsilon}{1-\lambda}$ then the tuple $\{x, t_x, \hat{f}_x\}$ is created or updated in the linked list F . The authors proved that the list F requires at most $\frac{r}{s-\epsilon}$ entries to store all of the items with decayed count greater than $\frac{s-\epsilon}{1-\lambda}$. Basically, λ -HCount can be considered a variant of Count-Min, designed to support frequent items detection in the time fading model.

The authors of λ -HCount proved that with a sketch requiring $\frac{e(1-\lambda) \ln(-\frac{M}{\ln p})}{\epsilon^2}$ space, where M is the number of distinct items and p is the success probability, and an additional data structure requiring at most $\frac{r}{s-\epsilon}$ space, where r is the number of hash functions used, their algorithm is able to estimate the frequent items decayed count with an error less than $\frac{\epsilon}{1-\lambda}$ with probability greater than p . The algorithm's analysis also shows that all of the items whose exact decayed count exceeds $\frac{s}{1-\lambda}$ will be output (there are no false negatives) and no items whose decayed count is less than $\frac{s-\epsilon}{1-\lambda}$ will be output.

The worst case complexity of λ -HCount depends on the time required for updating the sketch D and the linked list F . When an item is received from the stream, the algorithm computes r hash functions and updates r entries in D ; the linked list F is also updated accordingly. Since the linked list F is accessed through a hash function, and its update is done in constant time, overall the worst case complexity of per item update is $O(1)$. Therefore, the whole algorithm has worst case complexity $O(r)$, i.e., $O(\ln(-\frac{M}{\ln p}))$. The space complexity is given by the memory required by the sketch D and the linked list F . Overall, the worst case space complexity is $O(r \cdot m) = O(\frac{\ln(-\frac{M}{\ln p})}{\epsilon^2})$.

Algorithm 2 λ -HCount algorithm: the update phase

Require: λ : fading factor; ϵ : error bound; s : support threshold; x : received item; t : arrival time;

Ensure: update of sketch D and linked list F related to item x

```

1: procedure  $\lambda$ -HCount Update( $\lambda, \epsilon, s, x, t$ )
2:    $\hat{f}_x \leftarrow \infty$ 
3:   for  $k = 1$  to  $r$  do
4:      $y \leftarrow h_k(x)$ 
5:      $D[k, y].s \leftarrow D[k, y].s \cdot \lambda^{t-D[k, y].t} + 1$ 
6:      $D[k, y].t \leftarrow t$ 
7:     if  $D[k, y].s < \hat{f}_x$  then
8:        $\hat{f}_x \leftarrow D[k, y].s$ 
9:     end if
10:  end for
11:  if  $\hat{f}_x > \frac{s-\epsilon}{1-\lambda}$  then
12:    if  $x \in F$  then
13:      change its entry to  $\{x, \hat{f}_x, t\}$  and move it to the tail of the queue  $F$ 
14:    else
15:      if the queue  $F$  is full then
16:        delete the item at the head of the queue  $F$ 
17:      end if
18:      insert  $\{x, \hat{f}_x, t\}$  at the tail of the queue  $F$ 
19:    end if
20:  end if
21: end procedure

```

4. The FDCMSS algorithm

In this section, we introduce our algorithm, distinguishing three different phases: initialization, stream processing, and querying.

The algorithm's initialization, shown in pseudo-code as Algorithm 3, requires as input parameters ϵ , the error; δ , the probability of failure; ϕ , the support threshold; and t_{init} , a timestamp. Initialization returns a sketch D . The procedure starts deriving $d = \lceil \ln 1/\delta \rceil$, the number of rows in the sketch and $w = \lceil \frac{e}{2\epsilon} \rceil$, the number of columns in the sketch. We shall explain the reason why we set w to this value in Section 5.

Then, for each of the $d \times w$ cells available in the sketch D we allocate a data structure S with two Space Saving counters c_1 and c_2 . Given a counter $c_j, j = 1, 2$, we denote by $c_j.i$ and $c_j.f$ respectively the counter's item and estimated decayed count. Finally, we set the support threshold to ϕ , select d pairwise independent hash functions $h_1, \dots, h_d : [m] \rightarrow [w]$ mapping m distinct items into w cells, initialize the *count* variable, representing the total decayed count of all of the items in the stream (see Definition 4) to zero and the L variable (our landmark time) to t_{init} , which is a timestamp less than or equal to all of the items' timestamps. The worst case complexity of the initialization procedure is $O(\frac{1}{\epsilon} \ln \frac{1}{\delta})$.

Updating the sketch upon arrival of a stream item i with timestamp t_i , shown in pseudo-code as Algorithm 4, requires computing x , which is the forward decayed weight of the item, and incrementing *count* by x . Note that when computing x , we do not normalize the result (dividing by $g(t - L)$ where t is the query time, since we do not know in advance the query time); normalization occurs instead at query time. Then, we update the d cells in which the item is mapped to by the corresponding hash functions by using the Space Saving item update procedure. The worst case complexity of the update procedure is $O(\ln \frac{1}{\delta})$.

Finally, in order to retrieve the frequent items, a query can be posed when needed. Let t be the query time. The query, shown in pseudo-code as Algorithm 5, initializes R , an empty set, and then it inspects each of the $d \times w$ cells in the sketch D . For a given cell, we determine c_m , the counter in the data structure S with maximum decayed count. We normalize the decayed count stored in c_m dividing by $g(t - L)$, and then compare this quantity with $\phi \frac{\text{count}}{g(t-L)}$. If the normalized decayed count is greater, we pose a point query for the item $c_m.i$, shown in pseudo-code as Algorithm 6. If p , the returned value, is greater than $\phi \frac{\text{count}}{g(t-L)}$, then we insert in R the pair $(c_m.i, p)$.

The point query for an item j returns its estimated decayed count. After initializing the *answer* variable to infinity, we inspect each of the d cells in which the item is mapped to by the corresponding hash functions, to determine the minimum decayed count of the item. In each cell, if the item is stored by one of the Space Saving counters, we set *answer* to the minimum between *answer* and the corresponding counter's decayed count. Otherwise (none of the two counters monitors the item j), we set *answer* to the minimum between *answer* and the minimum decayed count stored in the counters. We return the normalized *answer*, dividing by $g(t - L)$.

From the previous discussion it is clear that our algorithm also solves the *decayed count estimation* problem for arbitrary items. Indeed, given an item, it suffices to pose a point query for that item. Finally, since the worst case complexity of a point query is $O(\ln \frac{1}{\delta})$, the worst case complexity of the query procedure is $O(\frac{1}{\epsilon} (\ln \frac{1}{\delta})^2)$. We shall argue in section 7 that a query only takes a few milliseconds and therefore its complexity is, in practice, negligible.

Algorithm 3 Initialize

Require: ϵ , error; δ , probability of failure; ϕ , threshold;

Ensure: a sketch $D[1 \dots d][1 \dots w]$ properly initialized

```

1: procedure INITIALIZE( $\epsilon, \delta, \phi, t_{init}$ )
2:    $d \leftarrow \lceil \ln 1/\delta \rceil$ 
3:    $w \leftarrow \lceil \frac{e}{2\epsilon} \rceil$ 
4:   for  $i = 1$  to  $d$  do
5:     for  $j = 1$  to  $w$  do                                 $\triangleright$  allocate a data structure  $S$  with two counters  $c_1, c_2$  for  $D[i][j]$ 
6:        $D[i][j] \leftarrow S$ 
7:     end for
8:   end for
9:   Set support threshold to  $\phi$ 
10:  Choose  $d$  pairwise independent hash functions  $h_1, \dots, h_d : [m] \rightarrow [w]$ 
11:   $count \leftarrow 0$ 
12:   $L \leftarrow t_{init}$                                         $\triangleright t_{init}$  must be  $\leq$  of all of the items' timestamps
13: end procedure

```

Algorithm 4 Process

Require: i , an item; t_i , timestamp of item i ;**Ensure:** update of sketch related to item i

```
1: procedure PROCESS( $i, t_i$ )                                 $\triangleright$  compute the decayed weight of item  $i$  and update the sketch
2:    $x \leftarrow g(t_i - L)$ 
3:    $count \leftarrow count + x$ 
4:   for  $j = 1$  to  $d$  do
5:      $S \leftarrow D[j][h_j(i)]$ 
6:     SPACESAVINGUPDATE( $S, i, x$ )
7:   end for
8: end procedure
```

Algorithm 5 Query

Require: t , query time**Ensure:** set of frequent items

```
1: procedure QUERY( $t$ )
2:    $R = \emptyset$ 
3:   for  $i = 1$  to  $d$  do
4:     for  $j = 1$  to  $w$  do
5:        $S \leftarrow D[i][j]$ 
6:       let  $c_1$  and  $c_2$  be the counters in  $S$ , and  $c_m$  the counter with maximum decayed count
7:        $c_m \leftarrow \text{ARGMAX}(c_1, c_2)$ 
8:       if  $\frac{c_m \cdot f}{g(t-L)} > \phi \frac{count}{g(t-L)}$  then
9:          $p \leftarrow \text{POINTESTIMATE}(c_m.i, t)$ 
10:        if  $p > \phi \frac{count}{g(t-L)}$  then
11:           $R \leftarrow R \cup \{(c_m.i, p)\}$ 
12:        end if
13:      end if
14:    end for
15:  end for
16:  return  $R$ 
17: end procedure
```

Algorithm 6 PointEstimate

Require: j , an item; t , query time

Ensure: estimation of item j decayed count;

```
1: procedure POINTESTIMATE( $j, t$ )
2:    $answer \leftarrow \infty$ 
3:   for  $i = 1$  to  $d$  do
4:      $\mathcal{S} \leftarrow D[i][h_i(j)]$ 
5:     let  $c_1$  and  $c_2$  be the counters in  $\mathcal{S}$ 
6:     if  $j == c_1.i$  then
7:        $answer \leftarrow \text{MIN}(answer, c_1.f)$ 
8:     else
9:       if  $j == c_2.i$  then
10:         $answer \leftarrow \text{MIN}(answer, c_2.f)$ 
11:      else
12:         $m \leftarrow \text{MIN}(c_1.f, c_2.f)$ 
13:         $answer \leftarrow \text{MIN}(answer, m)$ 
14:      end if
15:    end if
16:  end for
17:  return  $\frac{answer}{g(t-L)}$ 
18: end procedure
```

4.1. Example

Here, we provide a concise example of the FDCMSS update and query algorithms. Of course, illustrating a really significant example will require using a sketch of suitable dimensions, but, in the interest of clarity and to avoid wasting space, we use a sketch D with $d = 2$ rows and $w = 5$ columns. Setting d and w implicitly determine δ and ϵ . The other parameters are $\rho = 1.1$, $\phi = 0.025$ and $\lambda = 0.999$. Table 1 shows the state of the sketch after processing 1000 items coming from a stream derived from the universe $\mathcal{U} = \{x \in \mathbb{N} : 1 \leq x \leq 20\}$.

Table 2 depicts the sketch state after updating the data structure upon arrival of item 6 with timestamp 1001. The hash functions associated to each row map the item 6 respectively to column 5 in the first row and to column 3 in the second row of the sketch. Since item 6 was already monitored by a counter in sketch cells $D[1][5]$ and $D[2][3]$ (see Table 1), Space Saving increments these counters by adding the corresponding weight which is computed as $g(t - L) = (1/\lambda)^{t-L} = 2.72$ where $t = 1001$ is the timestamp of item 6 and $L = 0$ is the landmark time.

Upon arrival of item 5 with timestamp 1002 (see Table 3), the hash functions associated to each row map the item respectively to column 1 in the first row and to column 5 in the second row of the sketch. However, this time the item is not monitored and both counters are full in the sketch cells $D[1][1]$ and $D[2][5]$. Therefore, the counters monitoring the item with minimum weight are selected and updated by Space Saving evicting the monitored items, substituting them with item 5 and incrementing them by the corresponding weight.

We now show how to query the sketch to retrieve the frequent item candidates. We query the sketch after updating it upon arrival of item 5 and before processing the next item. When querying the sketch data structure, all of the involved weights are normalized dividing them by $(1/\lambda)^{t_q-L} = 2.725$, where $t_q = 1003$ is the query's timestamp.

The estimated normalized decayed count is equal at this point to $C = 632.671$; since $\phi = 0.025$, the threshold required for an item to be considered frequent is given by $\phi C = 15.817$.

We scan each of the sketch cells, determine the monitored item with maximum weight and, if the normalized weight of this item exceeds the threshold, we execute a point query for this item, which returns the normalized minimum weight associated to the item in the sketch. Then, we compare the normalized minimum weight to the threshold again, and consider the item frequent if it exceeds the threshold.

For instance, we determine that item 2 is frequent as follows. Since item 2 is the majority item in $D[1][1]$ and its normalized weight 203.78 is greater than the threshold, we execute a point query which produces the best frequency estimate for the item. In our case the point query for item 2 returns 198.08, and item 2 therefore is selected as

candidate frequent item because its normalized minimum weight is still greater than the threshold. Table 4 includes all of the frequent item candidates as returned by the query algorithm. None of the other items is selected as candidate frequent item. For instance, item 14 is not selected as candidate frequent item since its normalized minimum weight is $13.35 < 15.817$.

Table 1: Sketch state after 1000 updates

	Item	Weight	Item	Weight	Item	Weight	Item	Weight	Item	Weight
c_1	2	555.33	3	262.06	12	36.55	10	52.27	6	98.22
c_2	4	537.23	14	103.54	17	14.78	18	21.88	11	36.76
c_1	4	172.20	14	36.40	6	125.75	2	539.78	3	263.07
c_2	12	109.28	16	35.78	7	125.15	8	117.33	10	193.90

Table 2: Sketch state after arrival of item 6 with timestamp 1001

	Item	Weight	Item	Weight	Item	Weight	Item	Weight	Item	Weight
c_1	2	555.33	3	262.06	12	36.55	10	52.27	6	100.94
c_2	4	537.23	14	103.54	17	14.78	18	21.88	11	36.76
c_1	4	172.20	14	36.40	6	128.47	2	539.78	3	263.07
c_2	12	109.28	16	35.78	7	125.15	8	117.33	10	193.90

Table 3: Sketch state after arrival of item 5 with timestamp 1002

	Item	Weight	Item	Weight	Item	Weight	Item	Weight	Item	Weight
c_1	2	555.33	3	262.06	12	36.55	10	52.27	6	100.94
c_2	5	539.95	14	103.54	17	14.78	18	21.88	11	36.76
c_1	4	172.20	14	36.40	6	128.47	2	539.78	3	263.07
c_2	12	109.28	16	35.78	7	125.15	8	117.33	5	196.62

5. Error bound

In this section, we formally prove the error bound of our algorithm. Let f_i and \hat{f}_i be respectively the exact and estimated decayed count of item i . We denote by S a Space Saving summary, which is a data structure present in each one of the cells available in the sketch D . Let C be the decayed count of all of the items in the stream σ (see Definition 4). The main result of this section is the following theorem.

Theorem 1. $\forall k \in [m]$, \hat{f}_k estimates the exact decayed count f_k with error less than ϵC and probability greater than $1 - \delta$.

Proof. The algorithm is based on the use of d pairwise independent hash functions $h_1, \dots, h_d : [m] \rightarrow [w]$. In the following, we shall use related indicator random variables $I_{i,j,k}$ defined as follows:

Table 4: Frequent Item Candidates

Item	Normalized Decayed Weight
2	198.08
3	96.16
5	72.15
6	37.04

$$I_{i,j,k} = \begin{cases} 1 & h_i(k) = j \\ 0 & \text{otherwise} \end{cases}. \quad (6)$$

In other words, the indicator random variables $I_{i,j,k}$ are equal to one when the item $k \in [m]$ falls in the $D[i][j]$ cell. By pairwise independence of the hash functions, it follows that the expected value of $I_{i,j,k}$ is

$$\mathbb{E}[I_{i,j,k}] = \Pr[h_i(k) = j] = \frac{1}{w} = \frac{2\epsilon}{e}. \quad (7)$$

Denote by $S_{i,j}$ the sum of the decayed counts of the items falling in the cell $D[i][j]$, i.e.,

$$S_{i,j} = \sum_{k=1}^m f_k I_{i,j,k}. \quad (8)$$

Our algorithm processes each item falling in the same cell by using Space Saving with two counters. Now, we bound the error committed by Space Saving. By eq. (5), the error of an item monitored by a Space Saving counter is bounded by the sum of the decayed counts of the items which fall in the same cell divided by the number of counters. Therefore, denoting by $\hat{f}_{i,k}$ the estimated decayed count of item k returned by Space Saving for the $D[i][h_i(k)]$ cell and by f_k its exact decayed count, it holds that

$$\hat{f}_{i,k} - f_k \leq \frac{S_{i,j}}{2}. \quad (9)$$

By linearity of expectation,

$$\mathbb{E}[S_{i,j}] = \mathbb{E}\left[\sum_{k=1}^m f_k I_{i,j,k}\right] = \sum_{k=1}^m f_k \mathbb{E}[I_{i,j,k}] = \frac{C}{w} = \frac{2\epsilon C}{e}, \quad (10)$$

with C equal to the total decayed count, as defined in Definition 4. It holds that, on average,

$$\mathbb{E}[\hat{f}_{i,k} - f_k] \leq \frac{1}{2} \mathbb{E}[S_{i,j}] = \frac{2\epsilon C}{2e} = \frac{\epsilon C}{e}. \quad (11)$$

Since the error $\hat{f}_{i,k} - f_k$ is nonnegative, we can apply the Markov inequality, obtaining

$$\Pr[\hat{f}_{i,k} - f_k \geq \epsilon C] \leq \frac{\mathbb{E}[\hat{f}_{i,k} - f_k]}{\epsilon C} \leq \frac{\epsilon C}{e\epsilon C} = e^{-1}. \quad (12)$$

It follows that, when considering all of the d cells in which items are mapped to by the independent hash functions $h_i, i = 1, \dots, d$, and recalling that the estimated decayed count $\hat{f}_k = \min\{\hat{f}_{1,k}, \dots, \hat{f}_{d,k}\}$, we get

$$\begin{aligned} \Pr[\hat{f}_k - f_k \geq \epsilon C] &= \Pr[\min\{\hat{f}_{1,k}, \dots, \hat{f}_{d,k}\} - f_k \geq \epsilon C] \\ &= \Pr\left[\bigwedge_{i=1}^d (\hat{f}_{i,k} - f_k \geq \epsilon C)\right] \\ &= \prod_{i=1}^d \Pr[\hat{f}_{i,k} - f_k \geq \epsilon C] \\ &\leq e^{-d} = \delta. \end{aligned} \quad (13)$$

Consequently,

$$\Pr[\hat{f}_k - f_k < \epsilon C] > 1 - \delta. \quad (14)$$

□

Remark 1. Here, we explain the reason why we set $w = \frac{\epsilon}{2\epsilon}$ when we initialize the sketch data structure. In Theorem 1, letting w unspecified in equation (10), we can easily derive the relationship between d and w at the end of the Theorem: $d = \frac{\log \frac{1}{\epsilon}}{\log 2w\epsilon}$. Setting $w = \frac{\epsilon}{2\epsilon}$ we minimize the total space occupied by the sketch data structure. Indeed, the total space is $wd = w \frac{\log \frac{1}{\epsilon}}{\log 2w\epsilon}$; minimizing analytically wd with regard to w , we obtain the minimum for $w = \frac{\epsilon}{2\epsilon}$.

6. Correctness

We are going to formally prove the correctness of our algorithm. Before discussing its correctness, it is worth noting here that given a cell in the sketch D , the sum of the decayed counts stored by its two Space Saving counters is equal to the value that the Count-Min algorithm would store in that cell. However, Count-Min relies on an external heap data structure to keep track of frequent items. By using a data structure \mathcal{S} , with just two Space Saving counters per cell, we are able to dynamically maintain frequent items. Therefore, by using $O(dw)$ space as in Count-Min, our algorithm can solve both the *decayed count estimation* and the *approximate frequent items under forward decay* problems.

Here, we show that, with high probability, if an item is frequent, our algorithm will detect it. Indeed, given a cell, by using a data structure \mathcal{S} with two counters, we are able to detect the *majority item candidate* with regard to the sub-stream of items falling in that cell. Letting $S_{i,j}$ denote the total decayed count of the items falling in the cell $D[i][j]$, the majority item is, if it exists, the item whose decayed count is greater than $\frac{S_{i,j}}{2}$. The corresponding majority item candidate in the cell is the item monitored by the Space Saving counter whose estimated decayed count is maximum.

The main result of this section is the following theorem.

Theorem 2. *If an item i is frequent, then it appears as a majority item candidate in at least one of the d cells in which it falls with probability greater than or equal to $1 - (\frac{1}{2\phi w})^d$.*

Proof. Let k be a frequent item, $j = h_i(k)$, $D[i][j]$ one of the d cells in which the item k is mapped to by the corresponding hash function and $f_{min}^{(i,j)}$ the minimum of the two Space Saving counters available in the data structure \mathcal{S} monitoring the items falling in $D[i][j]$. Moreover, let $\hat{f}_{i,k}$ be the estimated decayed count of item k returned by Space Saving for the $D[i][h_i(k)]$ cell.

Our algorithm will not output the item k (and therefore will not be correct) iff for all of the d cells in which the item k is mapped to by the corresponding hash functions, the item k is not reported as a majority item candidate, so that

$$\hat{f}_{i,k} \leq f_{min}^{(i,j)}, \forall i = 1, \dots, d. \quad (15)$$

Indeed, by construction, our algorithm during a frequent item query only checks if an item is frequent when the item is reported in the cell as a majority item candidate. By assumption, since the item k is frequent, its decayed count is $f_k > \phi C$; since $f_k \leq \hat{f}_{i,k} \forall i = 1, \dots, d$, it holds that

$$\phi C < f_k \leq \hat{f}_{i,k} \forall i = 1, \dots, d. \quad (16)$$

Let $S_{i,j}$ denote the total decayed count of the items falling in the cell $D[i][j]$. By eq. (5), for the minimum decayed count in a cell $D[i][j]$ it holds that $f_{min}^{(i,j)} \leq \frac{S_{i,j}}{2}$.

We must now determine the probability that the event described in eq. (15) occurs. This is the probability of failing to correctly recognize a frequent item. Taking into account that $\phi C < \hat{f}_{i,k}$ and $f_{min}^{(i,j)} \leq \frac{S_{i,j}}{2}$, it holds that

$$S_{i,j} > 2\phi C. \quad (17)$$

By the previous argument, it follows that

$$\Pr[\hat{f}_{i,k} \leq f_{min}^{(i,j)}] < \Pr[S_{i,j} > 2\phi C]. \quad (18)$$

Reasoning as in section 5, $\mathbb{E}[S_{i,j}] = \frac{C}{w}$ with C equal to the total decayed count as defined in Definition 4. Using the Markov inequality we can bound the probability of failure (i.e., the probability of the item not being reported as a majority item candidate) in a single cell $D[i][j]$ taking into account eq. (17) and (18)

$$\Pr[\hat{f}_{i,k} \leq f_{min}^{(i,j)}] < \Pr[S_{i,j} > 2\phi C] \leq \frac{\mathbb{E}[S_{i,j}]}{2\phi C} = \frac{C}{2\phi w C} = \frac{1}{2\phi w}. \quad (19)$$

Therefore, we fail to identify a frequent item k when in all of the d cells $D[i][j]$, $i = 1, \dots, d$ in which the frequent item falls it is not reported as a majority item candidate in the data structure \mathcal{S} . We now estimate the corresponding probability. By eq. (19), the probability of failure is

$$\Pr\left[\bigwedge_{i=1}^d (\hat{f}_{i,k} \leq f_{min}^{(i,j)})\right] < \Pr\left[\bigwedge_{i=1}^d (S_{i,j} > 2\phi C)\right] \leq \left(\frac{1}{2\phi w}\right)^d. \quad (20)$$

Consequently, we succeed with probability greater than or equal to $1 - \left(\frac{1}{2\phi w}\right)^d$. \square

Remark 2. Since we proved in Theorem 2 that, if an item is frequent, then it appears as a majority item candidate in at least one of the cells in the sketch with high probability, it follows by the Space Saving design that two Space Saving counters are necessary and sufficient to determine this majority item candidate. Using more than two counters is useless for our purposes, and only wastes precious space.

7. Theoretical comparison of FDCMSS and λ -HCount

We provide here a thorough comparison of FDCMSS and λ -HCount. Both algorithms use a sketch data structure; FDCMSS is based on Count-Min and λ -HCount on the *hCount* algorithm [23]. However, λ -HCount relies on a backward decay exponential function, whilst FDCMSS can use either an exponential function or any other forward decay function. In particular, the use of a polynomial function allows more flexibility with regard to time fading. Indeed, using an exponential function the time fades faster, whilst with a polynomial function the times fades more slowly. Another advantage of using forward decay is that FDCMSS can easily deal with out of order arrival of stream items [9], something requiring significant effort to accommodate for λ -HCount.

Another difference is the use in λ -HCount of a dedicated data structure F to keep track of frequent items. Instead, our algorithm FDCMSS does not require additional space beyond its sketch data structure. Even though a query for frequent items requires in the worst case $O(\frac{1}{\epsilon}(\ln \frac{1}{\delta})^2)$, a query execution only takes a few milliseconds and therefore its complexity is, in practice, negligible (this has been verified in all of the experimental tests carried out). Moreover, queries are posed to FDCMSS from time to time whilst updates happen with high frequency, especially in high-speed streams. Therefore, FDCMSS has been designed to provide very fast updates, besides accurate results.

Theoretically, the main drawback of λ -HCount lies in the huge amount of space required to attain its error bound. In particular, the λ -HCount sketch requires $\frac{e(1-\lambda)\ln(-\frac{M}{\ln p})}{\epsilon^2}$ cells, where M is the number of distinct items and p is the success probability. Without taking into account the additional data structure F requiring $\frac{r}{s-\epsilon}$ entries, where r is the number of hash functions used, for a sketch using a total of $r \times m$ cells, we have:

$$r \times m = \left\lceil \frac{e(1-\lambda)\ln(-\frac{M}{\ln p})}{\epsilon^2} \right\rceil, \quad (21)$$

whilst, in our case, FDCMSS requires only

$$d \times w = \left\lceil \ln \frac{1}{\delta} \frac{e}{2\epsilon} \right\rceil. \quad (22)$$

As an example, fixing $\lambda = 0.99$, $M = 1048575$, $p = 0.96$ and $\epsilon = 0.001$, a total of $r \times m = 463779$ cells are required by λ -HCount. In order to achieve the same success probability p , in FDCMSS we need to set $\delta = 0.04$ and $\epsilon = 0.001$, so that $p = 1 - \delta = 0.96$ and a total of just $d \times w = 4375$ cells is required instead by our algorithm.

Let us now consider only the sketch size, without taking into account the λ -HCount additional data structure F required for tracking the frequent items. Each cell in the λ -HCount sketch stores a decayed count (a double, 8 bytes) and a timestamp (a long, 8 bytes), whilst a FDCMSS cell stores two Space Saving counters. A counter keeps track of an item (an unsigned int, 4 bytes) and its decayed count (a double, 8 bytes). Therefore, λ -HCount requires 16 bytes per cell and FDCMSS 24 bytes per cell.

Figures 1a and 1b plot, using a logarithmic scale, the sketch size in kilobytes required respectively as a function of the success probability p and as a function of ϵ for both algorithms. Here, we have fixed for the first plot the values $\lambda = 0.99$, $M = 1048575$, $\epsilon = 0.001$, and let p vary from 0.7 to 0.99. Similarly, for the second plot we have fixed $\lambda = 0.99$, $M = 1048575$, $p = 0.96$, and let ϵ vary from 0.001 to 0.01.

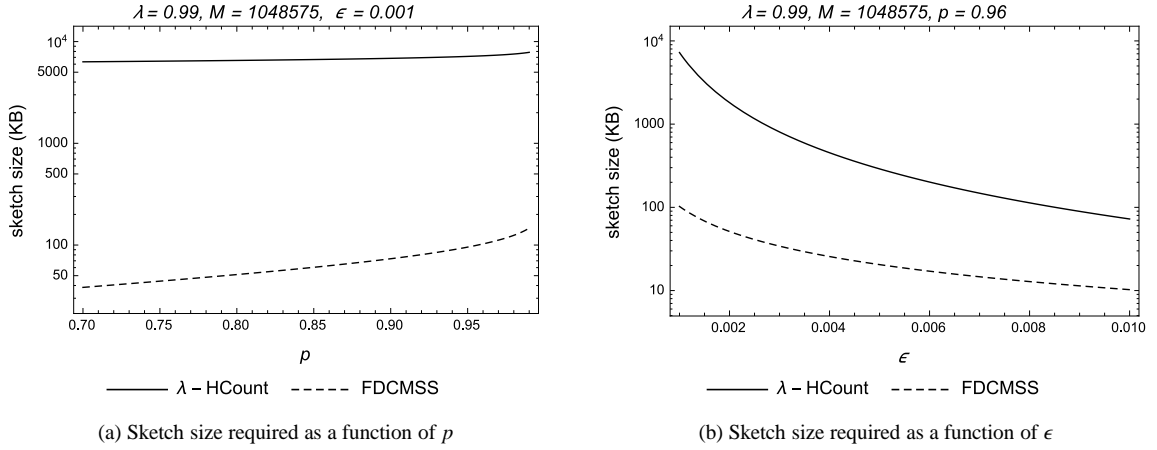


Figure 1: Sketch sizes in Kilobytes

It is immediate verifying that, in order to attain its theoretical error bound, λ -HCount requires a huge amount of space. On the contrary, FDCMSS achieves its bound using a tiny fraction of the space required by λ -HCount.

8. Experimental results

In this section, we report experimental results on both synthetic and real datasets. Here, we thoroughly compare our algorithm against λ -HCount [7] with regard to the performances on synthetic and real datasets.

8.1. Synthetic datasets

We have implemented FDCMSS and λ -HCount in C++. FDCMSS uses the *xxhash* hash function, and λ -HCount the *FNV* hash function as stated by its authors in [7]. The code has been compiled using the clang c++ compiler v7.0 on Mac OS X v10.11.2 with the following flags: `-Os -std=c++11`. We recall here that, on Mac OS X, the optimization flag `-Os` provides better optimization than the `-O3` flag and is the standard for building the release build of an application. The tests have been carried out on a machine equipped with 16 GB of RAM and a 3.2 GHz quad-core Intel Core i5 processor with 6 MB of cache level 3.

Regarding synthetic datasets, the input distribution used in our experiments is the Zipf distribution. For each different value of n (number of items), ϕ (support threshold), ρ (skew of distribution) and sketch size, the algorithms have been run 20 times using a different seed for the pseudo-random number generator associated to the distribution (using the same seeds in the corresponding executions of different algorithms). For each input distribution generated, the results have been averaged over all of the runs. The input elements are 32 bits unsigned integers.

In order to provide a fair comparison of the algorithms, we make sure that the decayed frequencies computed by both algorithms are equal. To this end, we use in FDCMSS the same exponential decay function $g(t - L) = (\frac{1}{\lambda})^{t-L}$ (in which the landmark time is $L = 0$) and the same $\lambda = 0.99$ parameter as in λ -HCount. This way, for a given input stream, the decayed counts of the input items and the set of frequent items computed by an exact algorithm are the

same for both algorithms. However, it is worth reporting here that the use of a forward decay polynomial function provides even better results for FDCMSS in terms of speed, measured as *updates per millisecond*.

We compare our algorithm against λ -HCount taking into account the following standard metrics: *recall*, *precision*, *mean absolute error*, *max absolute error*, *96-th percentile absolute error*, *updates per millisecond*. For each metric, we plot the values (mean and confidence intervals) obtained varying n , ϕ , ρ and the sketch size. In particular, for each plot we shall always compare the algorithms by using exactly the same sketch size in kilobytes. We will not take into account the additional space required by λ -HCount for its F data structure.

Recall, shown in Figure 2, is the total number of true frequent items reported over the number of true frequent items given by an exact algorithm. Therefore, an algorithm is correct iff its recall is equal to 1 (or 100%). We note here that λ -HCount recall is always 1 since the algorithm inserts an item in the F data structure only when that item has been detected as frequent. FDCMSS may instead provide a recall value lower than 1 (this follows immediately by Theorem 2; however, the probability of failing to detect a frequent item may be made arbitrarily small and close to zero by the user, setting appropriately the input parameters δ and ϵ). This happened in our experiments in only one case, when using a very small sketch size of only 6 KB. However, the measured recall value is 99.58% even in this extreme case.

Precision, shown in Figure 3, is defined as the total number of true frequent items reported over the total number of items reported. As such, this metric quantifies the number of false positives outputted by an algorithm. It follows that, from this point of view, an algorithm's precision should ideally be 1 (or 100%). The precision achieved by FDCMSS is 1 in the majority of the experiments, and our algorithm outperformed λ -HCount in particular when varying n and the sketch size, whilst providing anyway higher precision when varying ϕ and ρ .

Denoting with f_i the true decayed count of item i and with \hat{f}_i the corresponding decayed count computed by an algorithm, then the absolute error is, by definition, the difference $|f_i - \hat{f}_i|$. Denoting with M the number of distinct items in a stream (i.e., the stream domain size), the *mean absolute error* is then defined as $\sum_{i=1}^M \frac{|f_i - \hat{f}_i|}{M}$, i.e., the mean of the absolute errors. Similarly, the *max absolute error* is defined as $\max_i |f_i - \hat{f}_i|$. Finally, consider the M absolute errors in ascending sorted order: the 96-th percentile is the absolute error found in the position corresponding to 96% of M .

Figures 4, 5 and 6 show, respectively, the mean, max and 96-th percentile of absolute errors. FDCMSS outperforms λ -HCount, in particular with regard to the experiments in which we vary n and ϕ .

Regarding the skew, when the value of ρ increases, a reduction in the error committed and a corresponding increase in the accuracy is expected, owing to the Zipfian distribution. Indeed, the number of frequent items depends on the value of ρ as follows. Increasing ρ decreases the number of distinct items, i.e., we have less items but with higher frequency. Vice versa, decreasing ρ , we have more items but with lower frequency. It follows that, considering the Count-Min sketch, the number of collisions decreases because there are less distinct numbers and the sketch cells can therefore better estimate the items' frequencies.

Let us now discuss the actual performances of the algorithms in terms of updates per millisecond, where an update is defined respectively as in Algorithm 2 for λ -HCount and Algorithm 4 for FDCMSS. As shown in Figure 7, FDCMSS outperforms λ -HCount in all of the experiments carried out.

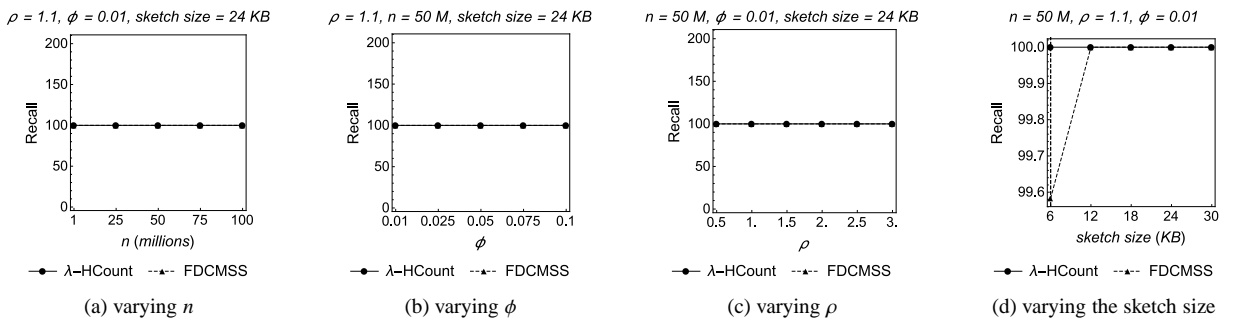


Figure 2: Recall (mean and confidence interval)

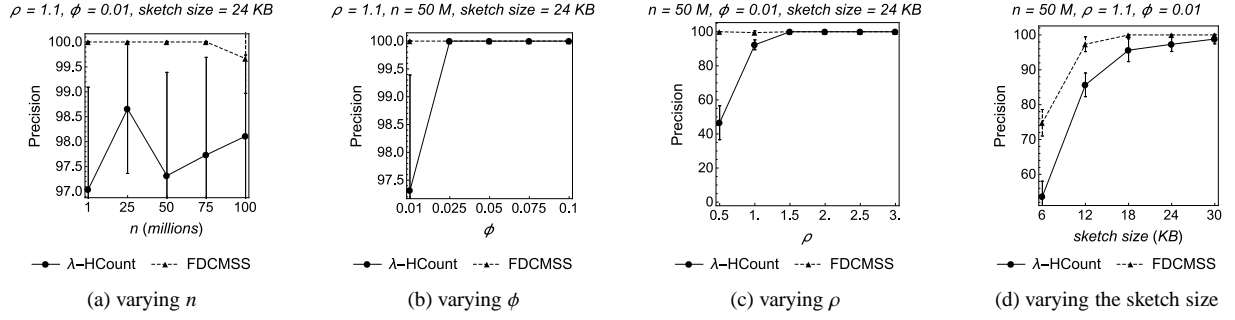


Figure 3: Precision (mean and confidence interval)

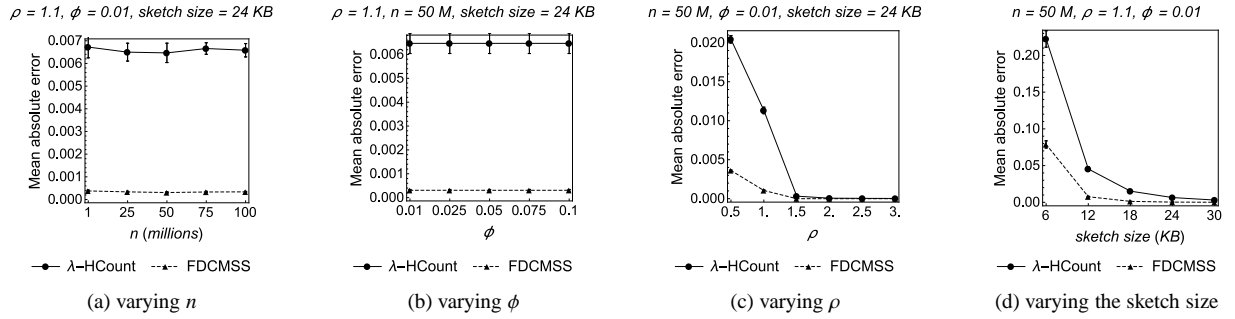


Figure 4: Mean absolute error (mean and confidence interval)

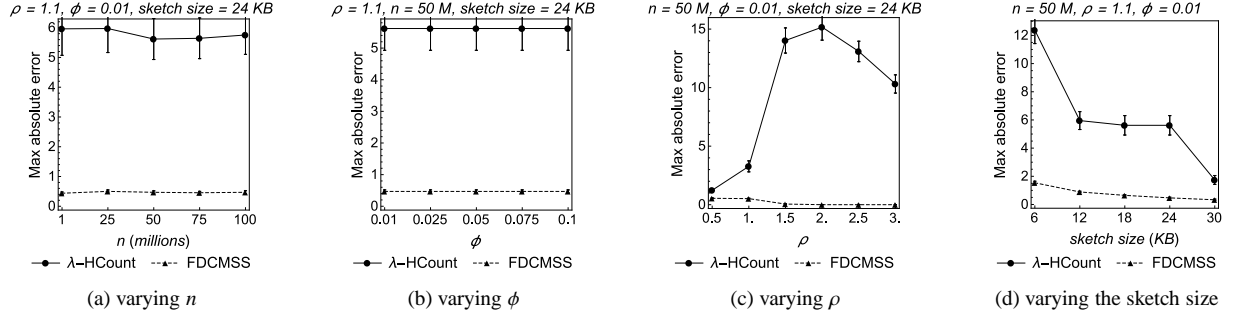


Figure 5: Max absolute error (mean and confidence interval)

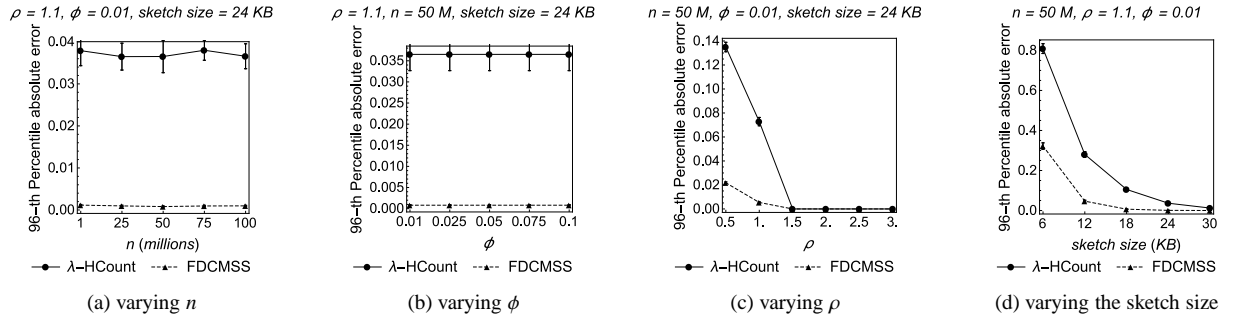


Figure 6: 96-th Percentile absolute error (mean and confidence interval)

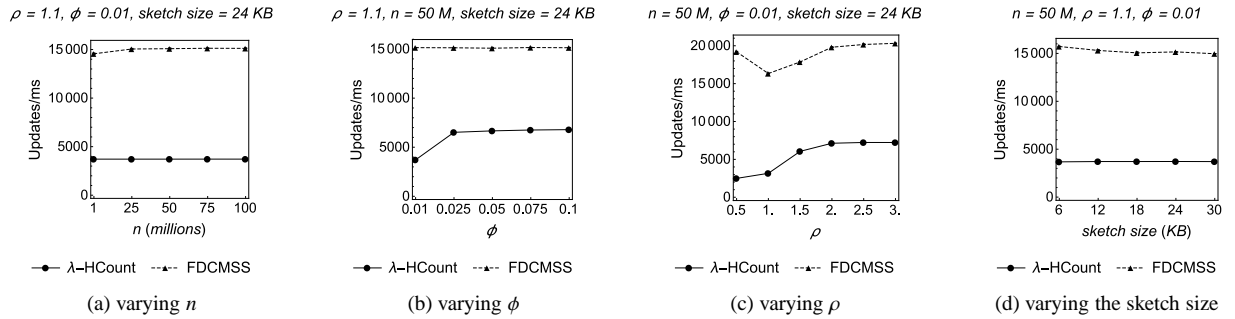


Figure 7: Updates/ms (mean and confidence interval)

Table 5: Statistical characteristics of the real datasets

	Kosarak	Retail	Q148	Webdocs
Count	8019015	908576	234954	299887139
Distinct items	41270	16470	11824	5267656
Min	1	0	0	1
Max	41270	16469	149464496	5267656
Mean	2387.2	3264.7	3392.9	122715
Median	640	1564	63	1988
Std. deviation	4308.5	4093.2	309782.5	549736
Skewness	3.5	1.5	478.1	6.1

8.2. Real datasets

We also provide experimental results for real datasets from several domains [1] [13]. These datasets are public domain and commonly utilized for data mining experiments. Moreover, a wide variety of statistical features, reported in Table 5, characterize the datasets, described below.

Kosarak: This is a click-stream dataset of a Hungarian online news portal. It has been anonymized, and consists of transactions, each of which is comprised of several integer items. In the experiments, we have considered every single item in serial order.

Retail: This dataset contains retail market basket data coming from an anonymous Belgian store. Again, we consider all of the items belonging to the dataset in serial order.

Q148: Derived from the KDD Cup 2000 data, compliments of Blue Martini, this dataset contains several data. The ones we use for our experiments are the values of the attribute "Request Processing Time Sum" (attribute number 148), coming from the "clicks" dataset. A pre-processing step was required, in order to obtain the final dataset. We had to replace all of the missing values (appearing as question marks) with the value of 0.

Webdocs: This dataset derives from a spidered collection of web html documents. The whole collection contains about 1.7 millions documents, mainly written in English, and its size is about 5 GB. The resulting dataset, after preliminary filtering and pre-processing, has a size of about 1.48 GB.

The experiments have been carried out by varying respectively ϕ and the sketch size. As we did for synthetic datasets, for each plot we always compare the algorithms by using exactly the same sketch size in kilobytes. We are not taking into account the additional space required by λ -HCount for its F data structure.

Figures 8, 11, 14 and 17 show precision and recall for the datasets under examination. In all of the experiments, FDCMSS and λ -HCount achieve 100% recall. Regarding precision, FDCMSS outperforms λ -HCount or provides the same precision. Indeed, both algorithms achieve 100% precision when varying ϕ , but FDCMSS outperforms λ -HCount when varying the sketch size.

Figures 9, 12, 15 and 18 show mean and max absolute errors. Clearly, FDCMSS outperforms λ -HCount in all of the experiments carried out. Figures 10, 13, 16 and 19 show the 96-th percentile of absolute error and the number of updates per millisecond. FDCMSS outperforms λ -HCount in all of the experiments carried out for these metrics.

9. Conclusions

We have presented the design and implementation of FDCMSS, a new algorithm for mining frequent items in the time fading model. Our algorithm is sketch based, and cleverly combines key ideas borrowed from forward decay, the Count-Min and the Space Saving algorithms. We have formally proved the correctness of our algorithm and shown, through extensive experimental results, that FDCMSS outperforms λ -HCount, a recently developed algorithm, with regard to speed, space used, precision attained and error committed on both synthetic and real datasets. Future work include parallelizing the algorithm on both shared-memory and message-passing architectures.

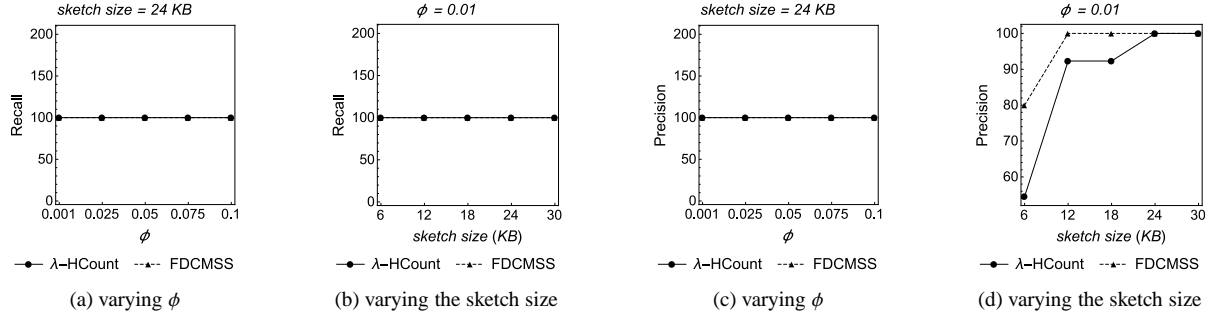


Figure 8: Kosarak: recall and precision (mean and confidence interval)

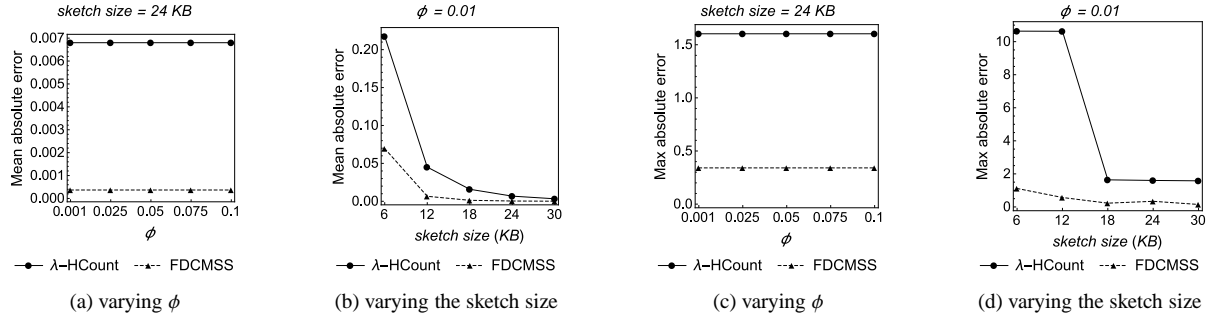


Figure 9: Kosarak: Mean and max absolute error (mean and confidence interval)

Acknowledgment

We are indebted to the unknown referees for enlightening observations, which helped us to improve the paper.

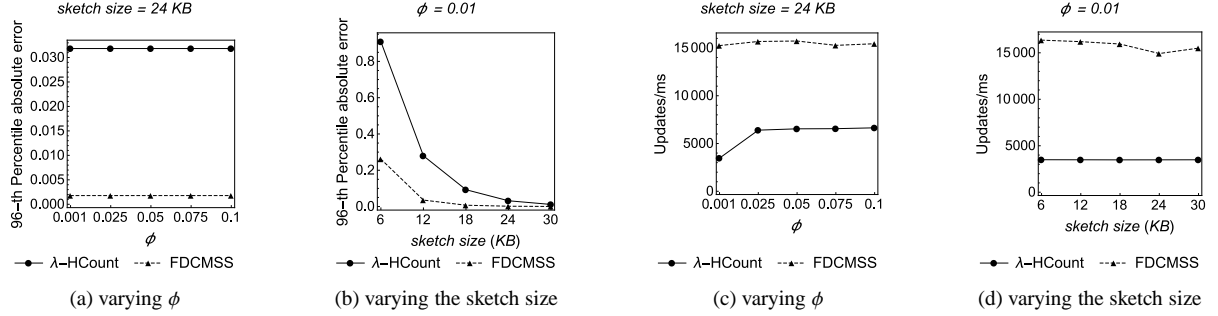


Figure 10: Kosarak: Percentile absolute error and updates/ms (mean and confidence interval)

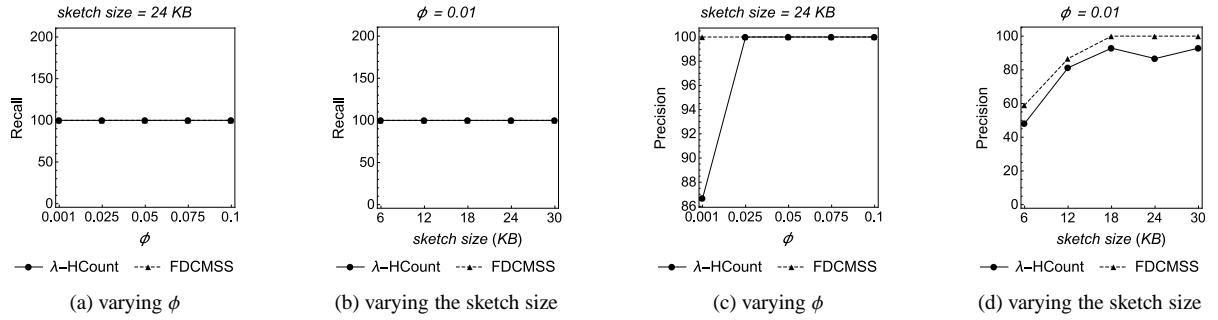


Figure 11: Retail: recall and precision (mean and confidence interval)

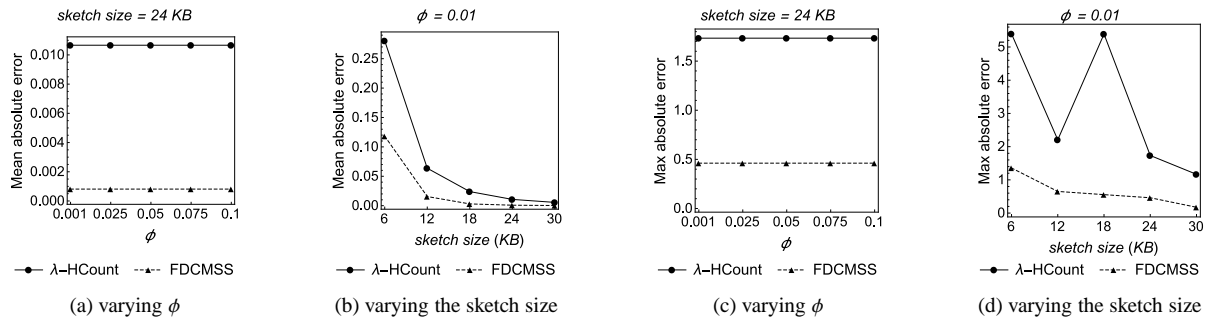


Figure 12: Retail: Mean and max absolute error (mean and confidence interval)

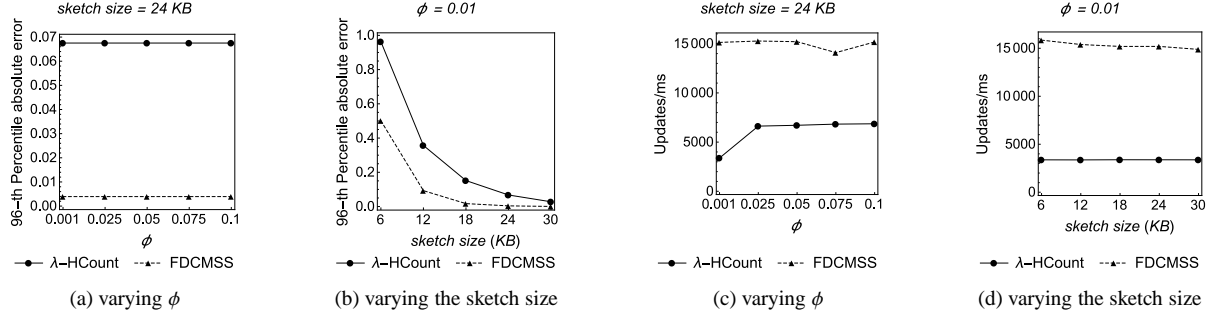


Figure 13: Retail: Percentile absolute error and updates/ms (mean and confidence interval)

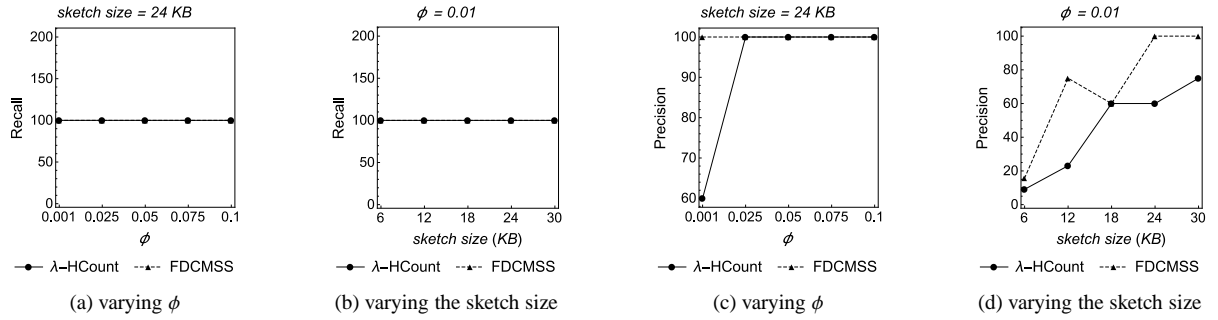


Figure 14: Q148: recall and precision (mean and confidence interval)

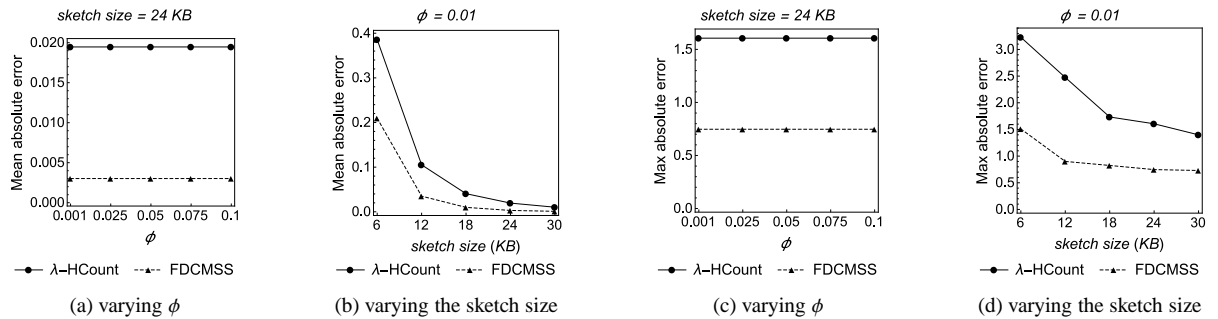


Figure 15: Q148: Mean and max absolute error (mean and confidence interval)

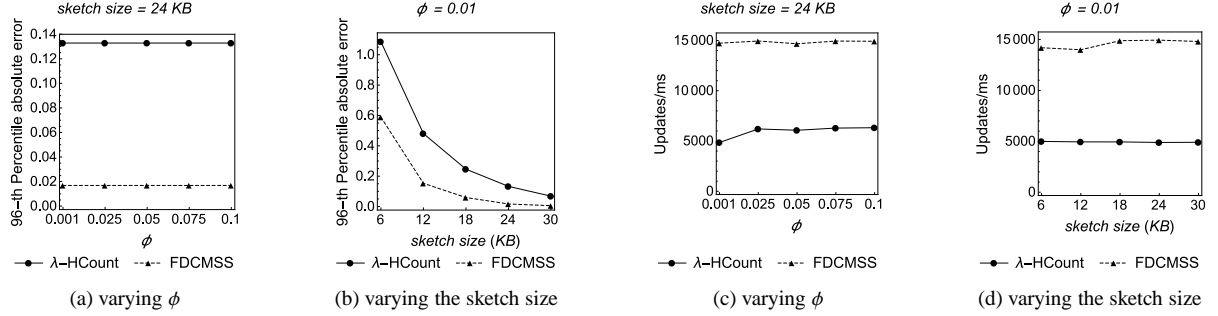


Figure 16: Q148: Percentile absolute error and updates/ms (mean and confidence interval)

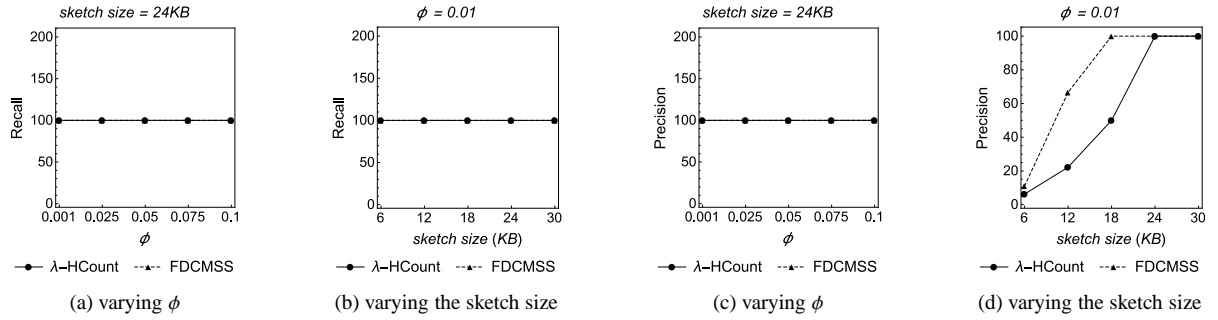


Figure 17: Webdocs: recall and precision (mean and confidence interval)

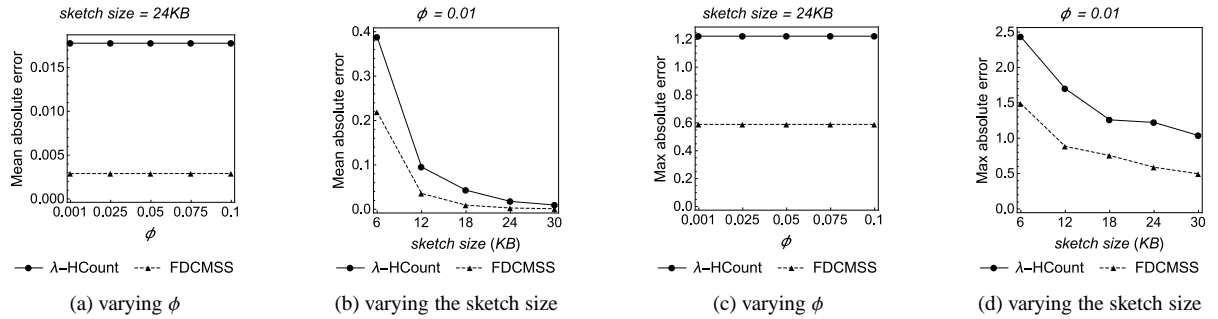


Figure 18: Webdocs: Mean and max absolute error (mean and confidence interval)

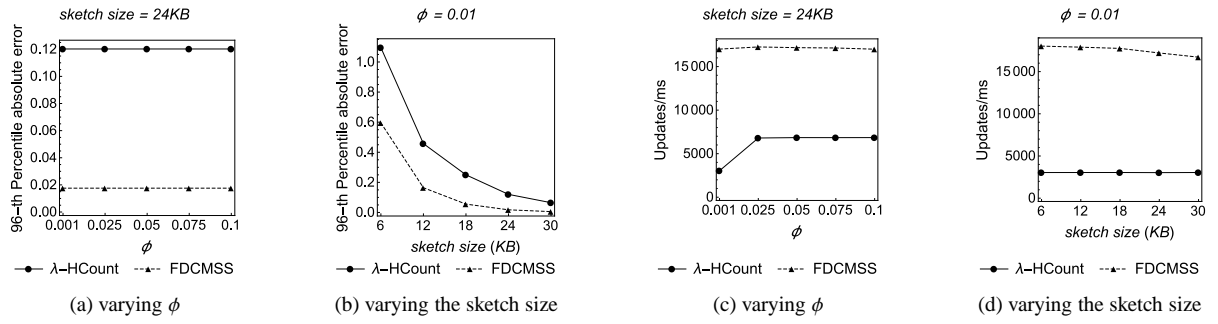


Figure 19: Webdocs: Percentile absolute error and updates/ms (mean and confidence interval)

References

- [1] *Frequent itemset mining dataset repository*, 2016 (accessed August 1, 2016). <http://fimi.ua.ac.be/data/>.
- [2] Kevin Beyer and Raghu Ramakrishnan, *Bottom-up computation of sparse and iceberg cubes*, Proceedings of the acm sigmod international conference on management of data. acm, new york, 1999, pp. 359–370.
- [3] Sergey Brin, Rajeev Motwani, Jeffrey D. Ullman, and Shalom Tsur, *Dynamic itemset counting and implication rules for market basket data*, Sigmod '97: Proceedings of the 1997 acm sigmod international conference on management of data, 1997, pp. 255–264.
- [4] Massimo Cafaro, Marco Pulimeno, and Piergiulio Tempesta, *A parallel space saving algorithm for frequent items and the hurwitz zeta distribution*, Information Sciences **329** (2016), 1–19.
- [5] Massimo Cafaro and Piergiulio Tempesta, *Finding frequent items in parallel*, Concurr. Comput. : Pract. Exper. **23** (October 2011), no. 15, 1774–1788.
- [6] Moses Charikar, Kevin Chen, and Martin Farach-Colton, *Finding frequent items in data streams*, Icalp '02: Proceedings of the 29th international colloquium on automata, languages and programming, 2002, pp. 693–703.
- [7] Ling Chen and Qingling Mei, *Mining frequent items in data stream using time fading model*, Information Sciences **257** (2014), 54–69.
- [8] G. Cormode, F. Korn, and S. Tirthapura, *Exponentially decayed aggregates on data streams*, Data engineering, 2008. icde 2008. ieee 24th international conference on, 2008April, pp. 1379–1381.
- [9] G. Cormode, V. Shkapenyuk, D. Srivastava, and Bojian Xu, *Forward decay: A practical time decay model for streaming systems*, Data engineering, 2009. icde '09. ieee 25th international conference on, 2009March, pp. 138–149.
- [10] Graham Cormode and Marios Hadjieleftheriou, *Finding the frequent items in streams of data*, Commun. ACM **52** (2009), no. 10, 97–105.
- [11] Graham Cormode and S. Muthukrishnan, *An improved data stream summary: the count-min sketch and its applications*, J. Algorithms **55** (2005), no. 1, 58–75.
- [12] ———, *What's hot and what's not: Tracking most frequent items dynamically*, ACM Trans. Database Syst. **30** (March 2005), no. 1, 249–278.
- [13] Michele Dallachiesa and Themis Palpanas, *Identifying streaming frequent items in ad hoc time windows*, Data Knowl. Eng. **87** (2013), 66–90.
- [14] Sudipto Das, Shyam Antony, Divyakant Agrawal, and Amr El Abbadi, *Thread cooperation in multicore architectures for frequency counting over multiple data streams*, Proc. VLDB Endow. **2** (August 2009), no. 1, 217–228.
- [15] Mayur Datar, Aristides Gionis, Piotr Indyk, and Rajeev Motwani, *Maintaining stream statistics over sliding windows: (extended abstract)*, Proceedings of the thirteenth annual acm-siam symposium on discrete algorithms, 2002, pp. 635–644.
- [16] Erik D. Demaine, Alejandro López-Ortiz, and J. Ian Munro, *Frequency estimation of internet packet streams with limited space*, Esa, 2002, pp. 348–360.
- [17] Ugo Erra and Bernardino Frola, *Frequent items mining acceleration exploiting fast parallel sorting on the {GPU}*, Procedia Computer Science **9** (2012), no. 0, 86–95. Proceedings of the International Conference on Computational Science, {ICCS} 2012.
- [18] Cristian Estan and George Varghese, *New directions in traffic measurement and accounting*, Imw '01: Proceedings of the 1st acm sigcomm workshop on internet measurement, 2001, pp. 75–80.
- [19] Min Fang, Narayanan Shivakumar, Hector Garcia-Molina, Rajeev Motwani, and Jeffrey D. Ullman, *Computing iceberg queries efficiently*, Proceedings of the 24th international conference on very large data bases, vldb. morgan-kaufmann, san mateo, calif., 1998, pp. 299–310.
- [20] Alexander Gelbukhl (ed.), *Computational linguistics and intelligent text processing, 7th international conference, cycling 2006*, Lecture Notes in Computer Science, vol. 3878, Springer-Verlag, 2006.
- [21] Phillip B. Gibbons and Yossi Matias, *Synopsis data structures for massive data sets*, Dimacs: Series in discrete mathematics and theoretical computer science: Special issue on external memory algorithms and visualization, vol. a., 1999, pp. 39–70.
- [22] Naga K. Govindaraju, Nikunj Raghuvanshi, and Dinesh Manocha, *Fast and approximate stream mining of quantiles and frequencies using graphics processors*, Proceedings of the 2005 acm sigmod international conference on management of data, 2005, pp. 611–622.
- [23] Cheqing Jin, Weining Qian, Chaofeng Sha, Jeffrey X. Yu, and Aoying Zhou, *Dynamically maintaining frequent items over a data stream*, In proc. of ckm, 2003, pp. 287–294.
- [24] Richard M. Karp, Scott Shenker, and Christos H. Papadimitriou, *A simple algorithm for finding frequent elements in streams and bags*, ACM Trans. Database Syst. **28** (2003), no. 1, 51–55.
- [25] Nishad Manerikar and Themis Palpanas, *Frequent items in streaming data: An experimental evaluation of the state-of-the-art*, Data Knowl. Eng. **68** (2009), no. 4, 415–430.
- [26] A. Manjhi, V. Shkapenyuk, K. Dhamdhere, and C. Olston, *Finding (recently) frequent items in distributed data streams*, Data engineering, 2005. icde 2005. proceedings. 21st international conference on, 2005April, pp. 767–778.
- [27] Gurmeet Singh Manku and Rajeev Motwani, *Approximate frequency counts over data streams*, In vldb, 2002, pp. 346–357.
- [28] Ahmed Metwally, Divyakant Agrawal, and Amr El Abbadi, *An integrated efficient solution for computing frequent and top-k elements in data streams*, ACM Trans. Database Syst. **31** (September 2006), no. 3, 1095–1133.
- [29] Jayadev Misra and David Gries, *Finding repeated elements*, Sci. Comput. Program. **2** (1982), no. 2, 143–152.

- [30] S. Muthukrishnan, *Data streams: Algorithms and applications*, Foundations and Trends in Theoretical Computer Science **1** (2005), no. 2, 117–236.
- [31] Rong Pan, Lee Breslau, Balaji Prabhakar, and Scott Shenker, *Approximate fairness through differential dropping*, SIGCOMM Comput. Commun. Rev. **33** (2003), no. 2, 23–39.
- [32] Pratanu Roy, Jens Teubner, and Gustavo Alonso, *Efficient frequent item counting in multi-core hardware*, Proceedings of the 18th acm sigkdd international conference on knowledge discovery and data mining, 2012, pp. 1451–1459.
- [33] Kanat Tangwongsan, Srikanta Tirthapura, and Kun-Lung Wu, *Parallel streaming frequency-based aggregates*, Proceedings of the 26th acm symposium on parallelism in algorithms and architectures, 2014, pp. 236–245.
- [34] Yu Zhang, *Parallelizing the weighted lossy counting algorithm in high-speed network monitoring*, Instrumentation, measurement, computer, communication and control (imccc), second international conference on, 2012, pp. 757–761.
- [35] Yu Zhang, Yue Sun, Jianzhong Zhang, Jingdong Xu, and Ying Wu, *An efficient framework for parallel and continuous frequent item monitoring*, Concurrency and Computation: Practice and Experience **26** (2014), no. 18, 2856–2879.